When should I use the FIND_FIRST_EX_LARGE_FETCH flag to FindFirstFileEx?

devblogs.microsoft.com/oldnewthing/20131024-00

October 24, 2013

Raymond Chen

Windows 7 introduces a new flag to the `FindFirstFileEx` function called `FIND_FIRST_EX_LARGE_FETCH`. The documentation says that it "uses a larger buffer for directory queries, which can increase performance of the find operation." This is classic MSDN-style normative documentation: It provides "just the facts". Far be it for MSDN to tell you how to write your application; the job of function-level documentation is to document the function. If you want advice, go see a therapist. If the reason why you're calling `FindFirstFileEx` is to enumerate through the entire directory and look at every entry, then a large buffer is a good thing because it reduces the number of round trips to the underlying medium. If the underlying medium is a network drive halfway around the world, the latency will be high, and reducing the number of calls reduces the overall cost of communication. Another case where you have high latency is if you are enumerating from an optical drive, since those tend to be slow to cough up data, and once you get the medium spinning, you want to get all the information you can before the drive spins the medium back down. On the other hand, if your underlying medium has low latency, then there isn't much benefit to using a large buffer, and it can be a detriment if the channel is low bandwidth, because transferring that large buffer will take a long time, which can result in long pauses on your UI thread. But what if you aren't enumerating with the purpose of reading the entire contents but rather are going to abandon the enumeration once you get the answer to your question? For example, maybe your function wants to enumerate the directory to see if it contains more than ten files. Once the tenth call to `FindNextFile` succeeds, you're going to abandon the enumeration. In this case, a large buffer means that the underlying medium is going to do work that you will end up throwing away.

Here's the above discussion summarized in a table, since people seem to like tables so much.

Scenario		**Use `FIND_FIRST_EX_LARGE_FETCH`?**
Enumerating entire directory	on UI thread	No[1]

	on background thread	Yes
Abandoning enumeration prematurely		No

[1] Actually, if you’re on a UI thread, you should try to avoid any directory enumeration at all.

Raymond Chen

Follow